NWG/RFC# 596 EAT3 6-DEC-73 01:30 20812
Second Thoughts on Telnet Go-Ahead TYPED ONLINE: Judged 12/12/73
FILENAME: Ga-mis notes fall 20812
PROOFED: MM J 4/18/74

Network Working Group Edward Taft
Request for Comments PARC-MAXC
Ref: Telnet Protocol Specification (NIC #15372) 1

SECOND THOUGHTS ON TELNET GO-AHEAD 2

INTRODUCTION 3

In this RFC we present objections to the requirement that hosts implement the Telnet Go-Ahead (GA) command, as specified in the Telnet Protocol Specification (NIC #15372). The thrust of these objections is in three major directions: 3a

1. The GA mechanism is esthetically unappealing, both to myself and to many other people I have talked to. I shall attempt to describe why this is so. 3a1

2. As specified in the Protocol, GA will not, in general, work; i.e. it will not serve its intended purpose unless nosts make various unwarranted assumptions about how other hosts operate. 3a2

3. GA is impossible for most hosts to implement correctly in all cases. This is certainly true of the PDP-10 operating systems with which I am familiar (10/50 and Tenex). 3a3

The purpose of this RFC is to advocate either complete removal of the GA mechanism or relegating it to the status of a negotiated option whose default state is that it be suppressed. 3b

TERMINOLOGY 4

"Half-duplex" is a two-way communication discipline in which transmission takes place in only one direction at a time and the receiving party is constrained not to transmit until the transmitting party has explicitly given up control of the communication path ("turned the line around"). 4a

This definition is distinct from a common (but incorrect) use of the terms "half-duplex" and "full-duplex" to designate local and remote character echoing. 4a1

"Reverse break" is a means by which a computer connected to a terminal by a half-duplex path may regain control of the path for further typeout after previously having relinquished t. 4b

This is the complement of the "break" or "attention" mechanism, implemented by all half-duplex terminals, by means of which the user may gain control of the line while it is in use by the computer. 4b1

ESTHETIC OBJECTIONS TO GA 5

One assumption that permeates the Telnet Protocol specification (and is explicitly stated on Page 7) is that the "normal" mode of communication between computers and terminals is half-duplex, line-at-a-time. While historically this is partially true, it is also clear, both within the ARPA Network community and elsewhere, that the trend is toward highly interactive man-machine communication systems which are difficult to implement under half-duplex communication disciplines. 5a

The GA mechanism is an attempt to solve a specific problem, that of switching control between computer and user in a subset of those hosts utilizing IBM 2741 or equivalent terminals. I say "a subset" because in fact the problem arises only in the case of 2741s lacking the reverse break capability. Many people log into TIPs from 2741s (with reverse break); from what experience I have had, I think the TIP does a very good job of turning the line around at the right moments. (I am told this is also the case at Multics). 5b

Given the trend toward more interactive communication, and given the fact that the terminals on the Network requiring a Go-Ahead mechanism are a distinct minority of all terminals, I think we should be reluctant to burden our protocols with kludges that are so clearly a concession to obsolete design. 5c

I have little doubt that before long somebody (if not IBM) will produce a full-duplex 2741-like terminal (indeed, perhaps it has already been done). There is an obvious need for a terminal with selectric quality keyboard and hard-copy better suited to interactive applications (i.e. full-duplex). 5c1

As a more practical consideration, it makes little sense to have the default state of the GA option be the one that benefits the least number of hosts and terminals. 5d

There is no question that most parties to Telnet communication will immediately negotiate to suppress GA. To do otherwise will double the amount of network traffic generated by character-at-a-time typein, and will increase it by a non-negligible amount even for line-at-a-time typein. 5d1

It strikes me as worthwhile to minimize the number of such "necessary" option negotiations, especially in view of the large number of TIPs and mini-hosts on the Network. Many such hosts must, due to resource constraints, implement only a limited subset of the available options. It follows, then, that the default state of all options should be the one most hosts will be willing to use. 5d2

WHY GA WON'T WORK 6

We now show that a server process's being "blocked on input" (as specified in the Protocol) is not in itself a sufficient condition for sending out GA. 6a

This is due to the fact that the user Telnet has no control over the packaging of a "line" of information sent to the server; rather, this is a function of the NCP, which must observe constraints such as allocation and buffering. Consider the following example: 6b

A user types a line of text, which is buffered by his host's user Telnet until he signals end-of-line. His keyboard then becomes locked (this being the behavior of half-duplex terminals while the computer has control of the line), and stays locked in anticipation of the server's eventual response and subsequent GA command. 6b1

The user Telnet transmits this text line over the connection; however, due to insufficient allocation or other conditions, the text actually gets packaged up and sent as two or more separate messages, which arrive at the server host in the correct order but separated by some amount of time. 6b2

The server Telnet passes the contents of the first message to the appropriate process, which reads this partial text line and immediately blocks for further input. At this moment (assuming the second message hasn't arrived yet), the server Telnet, in accordance with the Protocol, sends back a GA command. 6b3

The rest of the text then arrives; in response, the server process may generate a large volume of output. Meanwhile, however, the GA command has caused the user's keyboard to become unlocked and computer output thereby blocked. Hence we have a deadlock, which will be resolved only when the user recognizes what has happened and (manually) gives control back to the computer. 6b4

Of course, this particular problem is avoided if the Telnet protocol is modified to specify that the server Telnet will transmit GA only if the server process is blocked for input AND the most recent character passed to that process was end-of-line. 6c

I claim that this solution is bad in principle because it assumes too much knowledge on the part of the serving host as to what constitutes "end-of-line" in the using host. 6c1

Furthermore, the Protocol explicitly (and quite rightly) specifies that the user Telnet should provide some means by which a user may signal that all buffered text should be transmitted immediately, without its being terminated by end-of-line. 6c2

One must conclude, then, that in general the server Telnet has no precise way of knowing when it should send GA commands. 6d

IMPLEMENTATION PROBLEMS 7

The foregoing analysis illustrates the problems that arise with the GA mechanism in communication between servers and users whose normal mode of operation is half-duplex, line-at-a-time. When we turn to hosts that provide full-duplex service, such as the PDP-10s and many other hosts on the Network, the problems are much more severe. 7a

This is particularly true of operating system such as Tenex that exercise such tight control over terminal behavior that they prefer to operate in server echoing, character-at-a-time mode. This will probably become less necessary as protocols such as the Remote Controlled Transmission and Echoing Option come into general use, enabling servers to regulate echoing and break character classes in user Telnets. 7a1

Even in hosts such as 10/50 systems that provide reasonable service to line-at-a-time users for most subsystems (e.g. excluding DDT and TECO), GA is impossible to implement correctly. This is true for several reasons. 7b

First, there are a number of subsystems that never block for terminal input but rather poll for it or accept it on an interrupt basis. In the absence of typein, such processes go on to do other tasks, possibly generating terminal output. 7c

Processes of this sort come immediately to mind. The user Telnet, FTP, and RJE programs are implemented in this fashion by almost all hosts. 10/50 has a subsystem called OPSER, used to control multiple independent subjobs from a single terminal. 7c1

Since these programs never block for input, GA commands will never be sent by the server Telnet in such cases even though the processes are prepared to accept terminal input at any time. 7c2

Second, there is not necessarily a one-to-one relationship between processes and terminals, as seems to be assumed by the Telnet Protocol specification. 7d

For example, in Tenex one process may be blocked for terminal input while another process is generating output to the same terminal. (Such processes are typically parallel forks of the same job). 7d1

Third, there is the possibility of inter-terminal links, such as are provided in many systems. 7e

By this I do not mean special Telnet connections established between a pair of NVTs for the express purpose of terminal-to-terminal communication, as is suggested on page 9 of the Protocol specification. Rather, I am referring to facilities such as the Tenex LINK facility, in which any number and any mixture of local and Network terminals and processes may have their input and output streams linked together in arbitrarily complex ways. Clearly the GA mechanism will fall flat on its face in this case. 7e1

Also, the notion that one user of an inter-terminal link should have to "manually signal that it is time for a GA to be sent over the Telnet connection" in order to unlock another user's keyboard offends me to no end. 7e2

Finally, most systems provide means by which system personnel and processes may broadcast important messages to all terminals (e.g. SEND ALL in 10/50, NOTIFY in Tenex). Clearly such asynchronous messages will be blocked by a half-duplex terminal that has been irrevocably placed in the typein state by a previous GA. 7f

This strikes me as such an obvious problem that I am forced to wonder how half-duplex hosts handle it even for their local terminals. 7f1